DRAFT

Federated Community Cloud

Preliminary Findings Report

http://www.federatedcloudteam.org

April 30, 2012

PRELIMINARY DRAFT

Executive Summary

NIST is leading the development of a United States Government (USG) Cloud Computing Technology Roadmap. This roadmap will define and prioritize USG requirements for interoperability, portability, and security for cloud computing in order to support secure and effective USG adoption of Cloud Computing. The NIST Reference Architecture and Taxonomy Working Group is helping to advance to the USG Cloud Computing Technology Roadmap.

In January 2012, the Federated Community Cloud Team was formed to address Requirement #5 of NIST's Draft Technology Roadmap Volume I, which calls for "*frameworks to support seamless implementation of federated community cloud environments. (interoperability and portability guidance and technology).*"

The team met every Friday at 4 PM ET with rare exception to discuss the topic of community cloud federation. This report summarizes the work in progress from those discussions.

CONTENTS

PRELIMINARY DRAFT

Introduction

In the case in which a cloud deployment model is not implemented in one environment (e.g., private cloud or public) that accommodates the entire community of interest, there is a need to define and implement mechanisms to support the governance and processes that enable federation and interoperability between different cloud service provider environments. Such definitions and mechanisms allow combinations of resources to be assembled to form a general or mission-specific Federated Community Cloud.

The importance of such Community cloud approaches was clearly identified in the NIST-hosted Reference Architecture public working group. The reference architecture anticipated potential multi-cloud configurations such as Hybrid Cloud or those topologies involving a Cloud Broker.

Federation techniques have been applied across grids, data centers, and multinational collaborations to create extensive existing multiuser grid, cloud, platform-as-a-service and software-as-a-service offerings. Up to now, these have primarily been implemented in the context of infrastructures for closed communities, or for more general public access to computational and storage infrastructures operated for specific defined communities, such as those for the NIH, the Open Science Grid and US supercomputing Teragrid user communities, etc. To a lesser but still important extent, some corporations have implemented in-house computing infrastructures based on the same techniques.

To address questions, the Federated Community Cloud Team sought to address five Priority Action Plans (PAPs):

1. Define federated Community cloud requirements and scenarios.
2. Identify how Hybrid Cloud and Cloud Broker elements described in the cloud Reference Architecture can be leveraged and harmonized in federated Community Cloud settings.
3. Document current usage patterns and projected near-term trends in grid and cloud architectures with attention to tools used for effective support of federated user communities.
4. Present analysis of grid communities' applicability to federated cloud communities, including technology, trust infrastructure, & governance.
5. All stakeholders -- assess Intercloud efforts (e.g. SDO's) for applicability.

PRELIMINARY DRAFT

PAP 1A: Define federated Community cloud requirements.

The team discussed the need for a definition to facilitate a common understanding of Federated Community Cloud, and adopted the following for purposes of its discussions:

Federated Community Cloud. A cloud in which the resources are provisioned for exclusive use by a specific community of consumers from multiple organizations that have shared concerns.

Federated Community Cloud Capabilities

Based approaches sources from Identified Use Cases, Federated Community Cloud have a wide variation in requirements. In some cases, federations are implemented by dedicated network connections between clouds. In other cases, there are demands for robust security capabilities. Today, use cases can best ascribe requirements for Federated Community Clouds.

Taken from existing Use Cases applicable to federated clouds, we've compiled a list of relevant capabilities. It is not the intent to imply listed capabilities are hard requirements for any single use case. Along with other cloud commuting characteristics, Federated Cloud Computing may include one or more of the following capabilities:

- privacy
- security
- compliance adherence
- trust infrastructure
- common governance
- private communications

The following variations are included for reference and additional consideration:

- Federated Cloud: A Community Cloud where cloud RESOURCES ARE provisioned for exclusive use by a specific community of consumers from MULTIPLE organizations that have shared concerns. It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises.

  Typically, federated clouds will include common policies and/or technologies, and agreements that cover one or more of the following characteristics: mission, security requirements, policy, and compliance considerations, credentials, trust infrastructure and common governance

- Federated Cloud: A Community Cloud sharing common resources.
- Federated Cloud: a business model in which multiple external and internal cloud computing services deployed and managed to achieve a common goal.
- Federated Cloud: the deployment and management of external and internal cloud computing services to match business needs.

PRELIMINARY DRAFT

PAP 1B: Define federated Community cloud scenarios.

In order to facilitate discussions, the team determined that there was a need for examples of scenarios in which federated community clouds may prove especially beneficial or necessary. The team settled on two scenarios described below:

Scenario A: Catastrophic Event Response

Rapid configuration of dynamic and scalable response capability/capacity

Catastrophic events – whether caused by natural "acts of God"(Katrina), system failure (Bhopal), or hostile acts (911 attacks) – are characterized by disruption of both physical and human communication and logistical systems, resulting in loss of response capabilities, capacity, and scalability. Generally, both immediate and continuing response systems are de facto dependence on static information systems which are dependent on established physical infrastructure (including local wireless, physical data and computational systems, tethered networks, etc.). Various cloud services may provide the backup and scalable surge capability essential to continuity of operations of individual systems. However, effective event-level response often will require dynamic and multi-level federation of multiple cloud systems and resources, to effect a virtual, response-unique cross-cloud capability which is effective, responsive, and secure.

Example:

During the Katrina crises, families and other groups were separated not only from all resources but from each other. Both survivors and victims were found by many independent rescue efforts, including many without identification, health records, etc. Large groups, particularly of elderly or others with various chronic medical situations, were located in various rescue shelters. Families were separated, as well, resulting in unidentified children and others. There were no records or other repository by which survivors or bodies could be linked to others; and there were no pharmacy or other health records available to responders to assess individual medical or other needs.

In response to this event, an extensive capability was created linking information capabilities spanning a number of health provider organizations, government records (Medicare, Medicaid, etc.), and special capabilities of key large scale enterprises (including the San Diego Supercomputing Center via the Teragrid, Microsoft, etc.). Government officials at HHS also collaborated by creating legal methods for sharing of critical medical data, thus allowing the federation of certain parts of the resources. This ad hoc capability allowed for continuous identification and linking of families and groups; dynamic linking of critical medical information to patients and early responders (specifically allowing provision of time-critical medications); and the SDSC provided a multi-source dynamic linking capability which was used to link survivor identification, video, and oral interview data, resulting in the reunification of families, and other groups separated by the catastrophic events.

In the Japan earthquake/tsunami and resulting nuclear crises, the entire local data infrastructure was disrupted, and even the rescue/response venues were rendered dangerous in many cases. The initial, and much of the early continuing response had to operate with completely inserted virtual info-system infrastructure and capability. Similar challenges were faced during Bhopal and, in different ways, across the 911 venues. Once again, the challenges of such events requires the creation of virtual capabilities which in the future will require extensive federation of clouds, meshes, and varying requirements for security, assurance, accuracy and agile policy responsiveness.

In the future, we have both the need and the opportunity to dynamically configure heterogeneous cloud services, systems, and capabilities, in support of any disaster scenario. However, in order to meet the need for efficacy, scalability, and security; there must be standards which provide for the clarity and functionality essential: (a) at the

cloud edge; (b) for various types of cross-cloud application dynamic linkages; (c) for technical sharing of functions and of data across cloud boundaries; and, (d) for specification of security/assurance primitives. The key to success of these standards likely will lie in their definition more by functional/mission metrics, rather than by specific technical implementations. This approach will allow continuous innovation of the capability delivery, but clearly defined metrics at the point of provisioning and user dependence.

Scenario B: Specialized Distance Medical Care Emergency Trauma Response

Dynamically Reconfigurable Capability for Patient-centric, Integrated Trauma Medicine

Trauma medicine – particularly in the context of major event response situations – often requires considerable emergency consultative and delivery capability available only at major class 1 trauma centers. The need for distance (telemedicine) capability is a well-known challenge; the DoD and other enterprises have made considerable strides in provisioning significant capabilities. However, the need for dynamically reconfigurable capability is much more complex and emerges quickly in the following scenarios: (a) trauma centers are incorporated in the disruption zone (Katrina); (b) emergency care delivery is limited due to event challenges; and, (c) the trauma complexity is such that care requires dynamic configuration of consultative, data, health information, computational, and other types of services and resources which might in the future be served only via a virtual federation of heterogeneous cloud systems, infrastructures, etc.

The requirements for this scenario are functionally quite similar to those of the catastrophic event scenario. The key difference in this case is the patient-centric vs. event-centric nature of the virtual federated cloud mission. In the case of individualized trauma care, the federation standard must be built around the life critical, high medical information assurance and security, and quality of service requirements unique to a trauma situation. Thus, in this scenario, the efficacy metrics, for example, will focus much more on narrow performance and delivery capabilities, etc. The challenge of this scenario for federated clouds thus demonstrates the importance of functional standards at the edges, as well as exchange standards across the virtual infrastructures.

PRELIMINARY DRAFT

PAP 2: Identify how Hybrid Cloud and Cloud Broker elements described in the cloud Reference Architecture can be leveraged and harmonized in federated Community Cloud settings.

The team conceives that Cloud Brokers may evolve to facilitate federation for communities of trust.

The NIST definition of Cloud Broker indicates that brokers may provide service "intermediation" including value added services. Consider a Business Process as a Service (BPaaS) that automates the buying and management of cloud services. For example, it may automate purchases of services to scale applications without human interaction, and it may send alarms or discontinue purchases when established thresholds are reached.

It is then conceivable that the BPaaS may interface with platform services to establish communities of trust and to provide federation of services within those communities of trust. For example, the BPaaS may provide a single sign-on that interfaces with shared platform services, such as for role-based control or data transport.

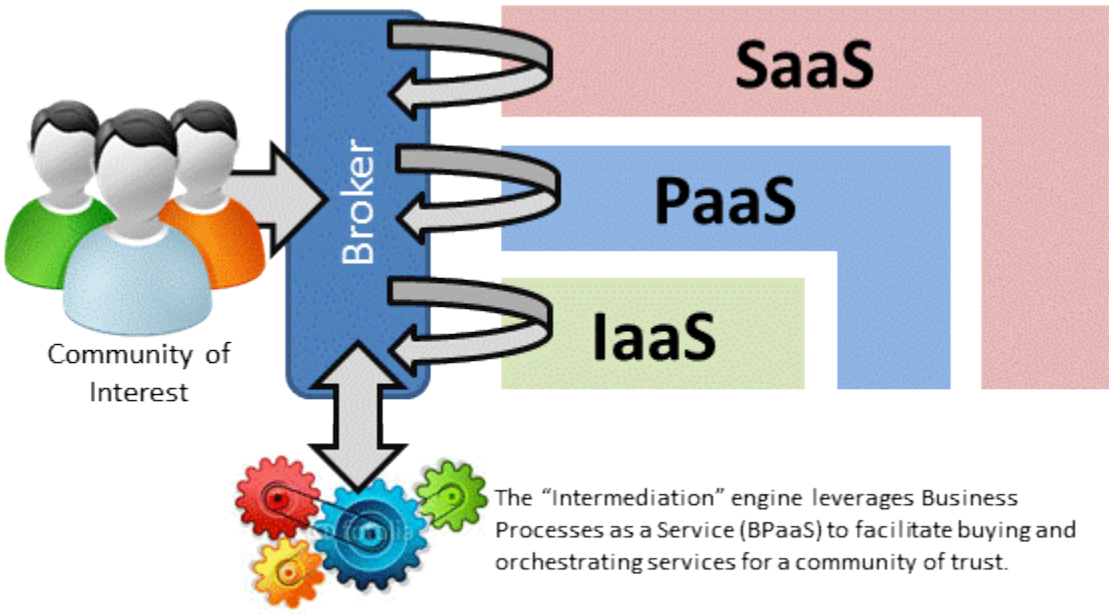

The Federated Community Cloud Team intends to further analyze possible ways that brokers can facilitate federation within the upcoming months. At this time, the team acknowledges the conceptual possibilities.

PRELIMINARY DRAFT

PAP 3: Document current usage patterns and projected near-term trends in grid and cloud architectures with attention to tools used for effective support of federated user communities.

The Federated Community Cloud Team acknowledged that the grid communities have already undertaken significant efforts to promote federation. The purpose of this PAP was to identify and raise awareness of such efforts so that they may be leveraged by the cloud communities.

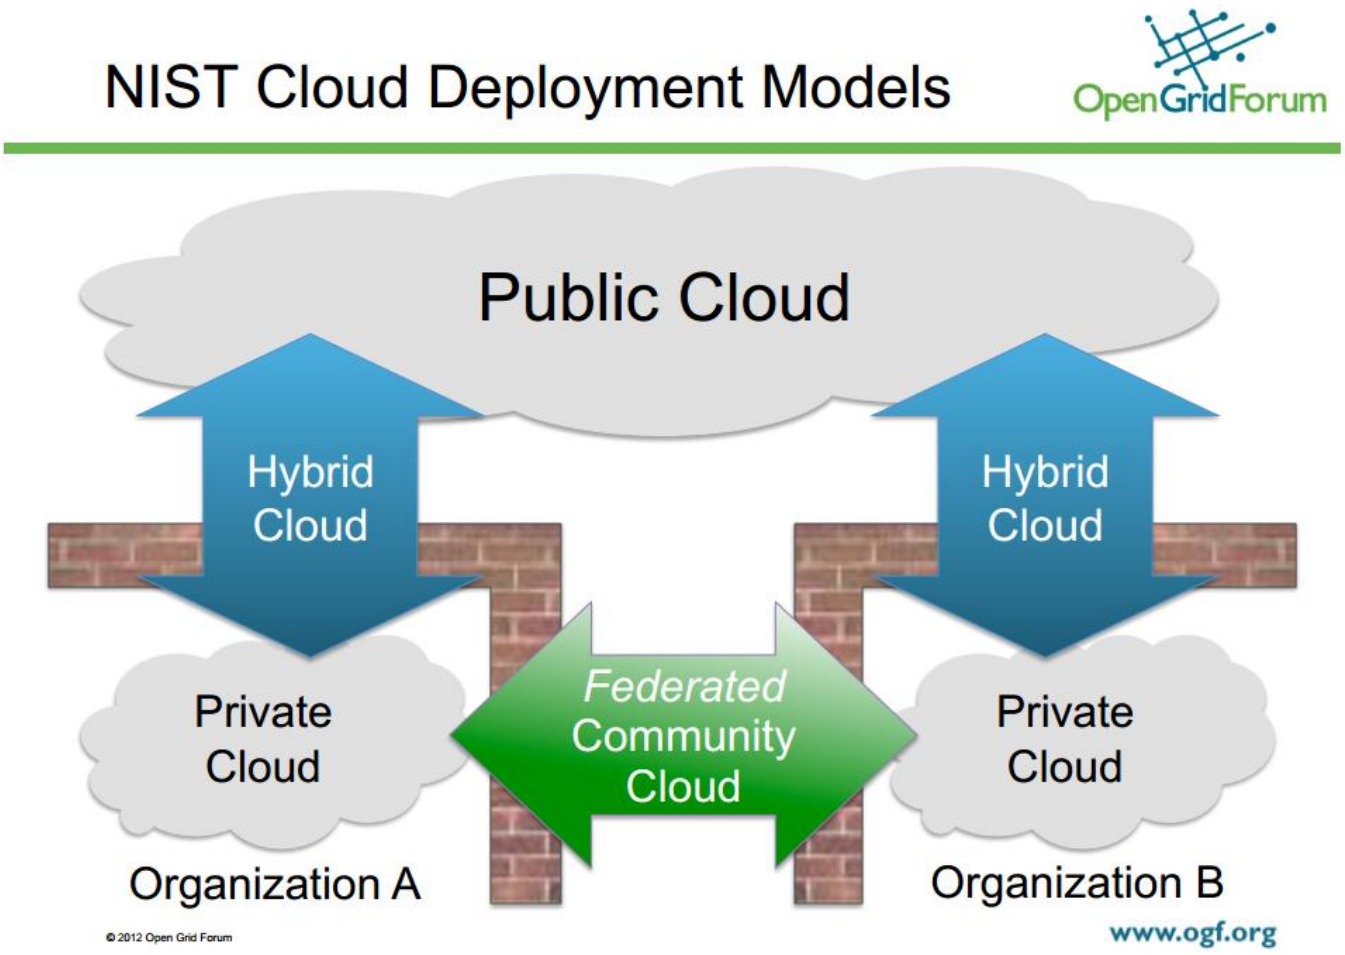

Figure 1: Open Grid Forum depiction of a Federated Community Cloud

The following are key findings:

A. Standards: The Open Grid Forum (OGF) has made available (verify) various applicable standards related to federated community grid and cloud computing, including:

- Federated Identity Management (FedSec-CG)
- Managing the Trust Eco-System (CA operations, AuthN/AuthZ)
- Virtual Organizations (VOMS)
- Job Submission and Workflow Management (JSDL, BES)
- Network Management (NSI, NML, NMC, NM)

PRELIMINARY DRAFT

- Secure, fast multi--party data transfer (GridFTP, SRM)
- Data Format Description (DFDL)
- Service Agreements (WS-Agreement, WS-Agreement Negotiation)
- Cloud Computing interfaces (OCCI)
- Distributed resource management (DRMAA, SAGA, etc.)
- Firewall Traversal (FiTP)

OGF continues to make new progress on basic core infrastructure specifications, as the following standards have been recently published:

- OCCI core, infrastructure and RESTful HTTP rendering (GFD. 183, 184 and 185 respectively).
- Data Management API within the GridRPC (GFD.186).
- OGSA-DMI Plain Web Service Rendering (GFD.187).
- WS-Iterator 1.0 (GFD.188)
- Relying Party Defined Namespace Constraints Policies in a Policy Bridge PKI Environment (GFD.189).
- Mapping between DFDL 1.0 Infoset and XML Data Model (GFD.190, supplements Data Format Description Language,GFD.174).
- Procedure for Registration of Subnamespace Identifiers in the URN:OGF Hierarchy (GFD.191).
- Web Services Agreement Specification (WS-Agreement) (GFD.192) and WS-Agreement Negotiation (GFD.193).
- Distributed Resource Management Application API Version 2 (DRMAA) (GFD.194, obsoletes 22, 130 & 133).

B. **Cooperative Agreements**: The Open Grid Forum (OGF) has established various cooperative agreements:
   - **OCCI and DMTF:** OGF published the the OCCI Core, Infrastructure and HTTP Rendering specifications as GFD.183, 184 and 185 respectively, and is working on a JSON rendering. We created a joint work register with DMTF and continue to follow their progress towards publishing the related CIMI specification.
   - **OGF and ISO:** OGF has been accepted as a Category A liaison with ISO JTC1 SC38 on Cloud Computing.
   - **OCCI and CDMI:** OGF has cooperative agreement w/SNIA and has held 4 jointly hosted Cloud Standards Plug-Fests so far; series continues.
   - **In process or planned:** OGF and TM Forum, OGF and CSA, OGF and IEEE, OGF and SIENA, NIST, GICTF, and others.

C. **European Grid Initiative (EGI) / Federated Cloud Task Force**

Copied from: https://wiki.egi.eu/wiki/Fedcloud-tf:FederatedCloudsTaskForce

EGI is a federation of national and domain specific resource infrastructure providers comprised of individual resource centres. Many of these resource centres have been experimenting with the deployment of virtualised management environments to improve the local delivery of services. Many of EGI's current and new user communities would like to access the flexibility provided by virtualisation across the infrastructure on demand in a 'cloud like' environment. Federating these individual virtualised resources is a major priority for EGI that has started with the EGI User Virtualisation Workshop, and the drafting of the EGI Cloud Integration Profile.

Objectives

PRELIMINARY DRAFT

- write a blueprint document for EGI Resource Providers that wish to securely federate and share their virtualised environments as part of the EGI production infrastructure;
- deploy a test bed to evaluate the integration of virtualised resources within the existing EGI production infrastructure for monitoring, accounting and information services;
- investigate and catalogue the requirements for community facing services based on or deployed through virtualised resources;
- provide feedback to relevant technology providers on their implementations and any changes needed for deployment into the production infrastructure;
- identify and work with user communities willing to be early adopters of the test bed infrastructure to help prioritise its future development;
- identify issues that need to be addressed by other areas of EGI (e.g. policy, operations, support & dissemination).

Activities

The Task Force mandate lasts eighteen months, from Sept 2011 to March 2013. The Task Force activities are organised in three, six-months long phases. During each phase, the Task Force evaluates a set of scenarios that an EGI federation of clouds should support. The scenarios are chosen by collecting use cases and requirements among user communities, resource providers and technology providers that have already adopted cloud computing or are planning to do so in a near future.

One or more work group is created inside the Task Force in order to evaluate each scenario. A leader is appointed for each work group and one or more collaborators are chosen among the Task Force members. The scenario evaluation performed by each work group is recorded in workbenchesand consists into:

- defining the set of capabilities that an EGI cloud infrastructure should have in order to support the given scenario;
- evaluating whether and what standards are available to implement the required capabilities;
- evaluating whether and what software solution is available to implement the required standards;
- evaluating the level of support of such software solutions by the Resource Providers members of the Task Force;
- evaluating the procedures to deploy the required software solutions into the infrastructure of the Resource Providers;
- devising tests to be run on the cloud infrastructure of the Resource Providers in order to simulate the given scenario.
-

PRELIMINARY DRAFT

Cloud Resources Status

The Task Force is developing a resource monitor solution for the clouds federation based on Nagios. Meanwhile, here a table showing the current status of the cloud resources made available by the resource providers that have joined the Task Force. This table is updated weekly by the resource providers.

Providers (■ = Available ■ = Not available)	User registration	User access	VM availability	Elastic IPs	Object Storage	Persistent Storage
CESGA (IBERgrid)						
CESNET (NGI CZ)	[1]	Sunstone OCCI v0.8 OCCI v1.1	VM suse (storage/16) with NET public (network/4)	No	Cumulus at carach3.ics.muni.cz:8888	GridFTP at carach4.ics.muni.cz:50000
CYFRONET (NGI PL)						
GWDG	[2]	■	■	■	■	■
FZ Jülich	Mail to Björn Hagemeier	• EC2 egi-cloud.zam.kfa-juelich.de:8773 • S3 egi-cloud.zam.kfa-juelich.de:3333	■	134.94.32.33 - 134.94.32.40	■	■
IGI						
IN2P3 (NGI FR)						
KTH	[3]	■	■		■	■
OerC (UK NGI)						
SARA (NGI NL)	[4]	■	■	■	■	■

PRELIMINARY DRAFT

PAP 4: Present analysis of grid communities' applicability to federated cloud communities, including technology, trust infrastructure, & governance.

The Grid community is involved in various initiatives to promote federation in support of Federated Identities and Delegation of Trust.

A. Federated Identity

Various Grid initiatives are underway to provide Federated Identity services for cloud communities.

Note that Federated Identity typically addresses Single Sign-On (Reuse of electronic identities), Delegation of Trust and Role-based authorization. Initiatives include:

- OpenID
- WebID
- Shibboleth
- Kerberos
- X.509 Certificates
- Simple Cloud Identity Model
- Windows Identity Foundation
- FIPS 201: Personal Identity Verification (PIV) of Federal Employees

B. Delegation of Trust

- Entity A allows Entity B to act on its behalf
- Entity B authorized to "impersonate" Entity A for specific purposes
- Implemented with X.509 Proxy Certificates (IETF RFC 3820)
- Example: Secure, Third-party Bulk Data Transfer – GridFTP (OGF GFD.20)

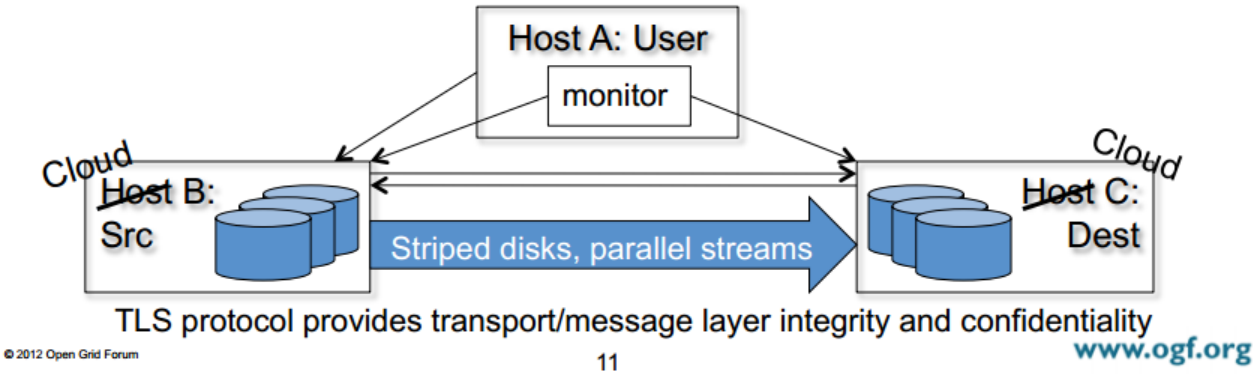

Figure 2: Delegation of Trust using TLS protocol

PRELIMINARY DRAFT

PAP 5: Assess Intercloud efforts (e.g. SDO's) for applicability to Federated Community Clouds.

This breakoff team was tasked with discovering capabilities and characteristics from Intercloud initiatives that may be applied to Federated Community Cloud infrastructures. The scope of this work is guided by the current definition of Community Cloud in document NIST SP800-145. Perspectives recognized are derived from relationships support by the characterizations of Cloud Brokers and Cloud Providers defined in document NIST SP500-252.

For the purposes of this work on Community Cloud Computing, "Intercloud Computing Initiatives" are described as work products created by SDOs and other organized computing industry associations specializing on "interconnecting cloud computing as a cloud of clouds".

Discussions surrounding concepts and capabilities of Intercloud computing started shortly after formation of the first cloud computing initiatives and intercloud initiatives are ongoing and maturing. Based on the works of current Intercloud initiatives, issues pertaining to Federated Community Cloud infrastructures may be classified in the following categories: discovery, engagement and operating between clouds.

The breakoff team identified a few groups that are producing works applicable to NIST's Federated Community Cloud infrastructures: the Alliance for Telecommunications Industry Standardization (ATIS), IEEE P2301 & P2302, Global Inter-Cloud Technology Forum (GICT) and the Open Grid Forum (OGF). Each group has taken independent approaches to work products, delivering four perspectives. Although independent initiatives, a common component present are use cases focusing their works.

Although not vetted though a formal process, we assume, US government Departments and Agencies will take on the roles of Cloud Broker, an Intermediary Cloud Provider and Cloud Provider. Segments of government will engage as Cloud Brokers and Cloud Providers, supplying services to Cloud Consumers originating from sectors including other areas of government, the general public, private sector, health care, public services, education, research, private sector, defense and foreign treaty partners to name a few.

In its course of business, we anticipated that US government Departments and Agencies will engage Cloud Brokers and Cloud Providers as suppliers. Adopting the role of Cloud Consumers, US government Departments and Agencies will share "common concerns" meeting the definition Cloud Consumer described in the SP 800-145 Community Cloud definition. It is also expected, the government will also act as a group of independent consumers, utilizing a single vendor's services as a traditional multi-tenant engagements and without additional requirements for consumer to consumer inter-operation.

If "financial advantage through purchase consolidation" is accepted as a "shared concern" which aligns to the definition of Community Cloud in document NIST SP800-145, the two latter use cases can be incorporated into a single "Community Cloud" Use Case.

Classifications of concerns and capabilities are work in progress.

The International Grid Trust Federation provides intercloud support:

- Provides trust accreditation among federation members
- Specifies how Certificate Authorities must be configured and operated
- AP/EU/TAG Policy Management Authorities verify compliance
- Self-audits with peer review, site visits, etc.

- IGTF members trust certificates signed by each other's CAs
- Scalability requires hierarchical trust relationships

http://www.igtf.net

PRELIMINARY DRAFT

Credits

The Federated Cloud Team included the following participants.

- Alan Sill
- Brett Berlin
- Bryan Ward
- Cary Landis
- Chris Braganza
- Chris Ferris
- Craig Lee
- Dave Harper
- Frederic de Vaulx
- Gary Mazzaferro
- Eugene Luster
- Gwendolyn Young
- JP Morgenthal
- James Ketner
- Jan Levine
- Karen Caraway
- Kazaz Harun
- Leslie Anderson
- Orit Levin
- Paul Sforza
- Shah Nawaz
- Shamun Mahmud
- Steve McGee
- Sundararajan Ramanathan
- Yin Lee
- Chris Uttenweiler
- Hussain Chinoy
- Kent Halloway
- Prabha Kumar
- Robert Marcus